

## METHOD AND SYSTEM FOR SCHEDULING ONLINE TARGETED CONTENT DELIVERY

### Background of the Invention

#### 5 Field of the Invention

The present invention relates generally to managing the delivery of content (such as advertisements) to online users and, more particularly, to a method and system for scheduling targeted content delivery.

#### Description of Related Art

10 Online advertising (typically banner advertisements) directed to Web users or subscribers has grown rapidly along with the growth of the Internet and e-commerce activity. Web advertising is frequently targeted to particular Web users. For example, it is known to profile Web users by determining their demographics and interests, and to selectively transmit advertisements to only  
15 those users having particular desired profiles. Information on users can be obtained from the users themselves through questionnaires or from other profiling techniques.

Online content delivery services typically contract with advertisers to deliver some specified number of advertisements (i.e., impressions) to Web users having  
20 particular profiles. In scheduling the delivery of advertisements, a number of variable factors are involved such as, e.g., when a given user will log on or for how long. Accordingly, the content delivery services typically pre-schedule all content deliveries for all users using very conservative estimates of scheduled deliveries in order to fulfill advertising contracts. Known content delivery systems are thereby  
25 very inefficient in utilizing inventory (available user screen real estate) and fulfilling advertiser contracts.

EXPRESS MAIL LABEL NO. EL 38491912005  
DATE OF DEPOSIT 1/23/01

A need exists for a method and system for more efficiently delivering targeted content to online users.

### Brief Summary of the Invention

The present invention is directed to a method and system for scheduling targeted content delivery to online users. In accordance with the invention, operations research and yield management techniques are utilized to forecast the availability of user screen real estate (i.e., inventory), optimize the use of surplus inventory, and generate on-demand selective content delivery schedules to transmit content to the users.

Briefly, the scheduler system in accordance with the invention reviews new advertising campaigns proposed by advertisers and predicts whether campaign objectives are achievable in view of forecasted inventory. The system generates a master delivery plan based on expected values to fulfill advertiser contracts and optimize usage of surplus inventory. The plan is periodically modified based on delivery feedback information. The system also dynamically generates individual user schedules on user login.

These and other features and advantages of the present invention will become readily apparent from the following detailed description wherein embodiments of the invention are shown and described by way of illustration of the best mode of the invention. As will be realized, the invention is capable of other and different embodiments and its several details may be capable of modifications in various respects, all without departing from the invention. Accordingly, the drawings and description are to be regarded as illustrative in nature and not in a restrictive or limiting sense with the scope of the application being indicated in the claims.

### **Brief Description of the Drawings**

For a fuller understanding of the nature and objects of the present invention, reference should be made to the following detailed description taken in connection with the accompanying drawings wherein:

FIGURE 1 is a schematic diagram illustrating of a representative network in which the inventive system is preferably implemented;

FIGURE 2 is a schematic diagram illustrating in greater detail the preferred architecture of the inventive system; and

FIGURE 3 is a flowchart illustrating in general the process for scheduling content delivery in accordance with the invention.

### Detailed Description of Preferred Embodiments

U.S. patent application Serial No. 09/558,755 filed April 21, 2000 and entitled "Method And System For Web User Profiling And Selective Content Delivery" is expressly incorporated by reference herein. That application discloses a method and system for profiling online users (who are also referred to herein as clients or subscribers) based on their observed surfing habits and for selectively delivering content, e.g., advertising, to the users based on their individual profiles.

The present invention is directed to optimizing content delivery to online users and can be implemented in a content delivery system such as, e.g., that disclosed in the above-mentioned application Serial No. 09/558,755. Very briefly, in accordance with the invention, operations research and yield management techniques are utilized to forecast the availability of user screen real estate (i.e., inventory), optimize the use of surplus real estate, and generate on-demand selective content delivery schedules to transmit content to the users. The content is displayed on user computer display devices, and can comprise advertising, e.g., in the form of banner advertisements or pop-up advertisements.

FIGURE 1 illustrates a general representative network in which the inventive scheduler system can be implemented. The network includes a plurality of client machines 10 operated by various individual users. The client machines 10 connect to multiple servers 12 via a communication channel 14, which is preferably the Internet. It may, however, alternatively comprise an Intranet or other known connections. In the case of the Internet, the servers 12 are Web servers that are selectively accessible by various clients. The Web servers 12 operate so-called "Web sites" and support files in the form of documents and pages. A network path to a Web site generated by the server is identified by a Uniform Resource Locator (URL).

One example of a client machine 10 is a personal computer such as a

Pentium-based desktop or notebook computer running a Windows operating system. A representative computer includes a computer processing unit, memory, a keyboard, a mouse and a display unit. The screen of the display unit is used to present a graphical user interface (GUI) for the user. The GUI is supported by the operating system and allows the user to use a point and click method of input, e.g., by moving the mouse pointer on the display screen to an icon representing a data object at a particular location on the screen and pressing on the mouse buttons to perform a user command or selection. Also, one or more "windows" may be opened up on the screen independently or concurrently as desired. The content delivered by the system to users is displayed on the screen.

Other types of client devices are also possible such as mobile Internet devices (e.g., Web-connected cellular phones and personal digital assistants).

Client machines 10 typically include browsers, which are known software tools used to access the servers 12 of the network. Representative browsers for personal computers include, among others, Netscape Navigator and Microsoft Internet Explorer. Client machines 10 usually access servers 12 through some private Internet service provider (ISP) such as, e.g., America Online. Illustrated in FIGURE 1 are multiple ISP "point-of-presence" (POP) systems, each of which includes an ISP POP server 16 linked to a group of client machines 10 for providing access to the Internet. Each POP server 16 is connected to a section of the ISP POP local area network (LAN) that contains the user-to-Internet traffic. The ISP POP server 16 can capture URL page requests from individual client machines 10 for use in user profiling (if user profiling is based on Web Surfing habits of users) and also to distribute targeted content to users. Also, as will be discussed in detail below, the inventive system also preferably includes a remote master server 18 linked to each ISP POP server 16 preferably through the Internet. The system software is

preferably distributed over the network at the ISP POP servers 16 and the master server 18 as will be discussed below.

As is well known, the World Wide Web is the Internet's multimedia information retrieval system. In particular, it is a collection of servers of the Internet that use the Hypertext Transfer Protocol (HTTP), which provides users access to files (which can be in different formats such as text, graphics, images, sound, video, etc.) using, e.g., a standard page description language known as Hypertext Markup Language (HTML). HTML provides basic document formatting and allows developers to specify links to other servers and files. These links include "hyperlinks," which are text phrases or graphic objects that conceal the address of a site on the Web.

A user of a client machine having an HTML-compatible browser (e.g., Netscape Navigator) can retrieve a Web page (namely, an HTML formatted document) of a Web site by specifying a link via the URL (e.g., www.yahoo.com/photography). Upon such specification, the client machine makes a transmission control protocol/Internet protocol (TCP/IP) request to the server identified in the link and receives the Web page in return.

FIGURE 2 illustrates in greater detail the preferred scheduler system architecture. (For simplicity, only a single POP server 16 and a single client 10 are shown.) As illustrated, the master server system 18 includes various software components for managing content delivery including a Dynamic Campaign Manager 50, a Capacity Forecaster 52, a Delivery Manager 54, an Inventory Manager 51, system configuration information 53, and a matcher 56. The master server system 18 also includes a master database 60 storing advertisements and user profiles. Also, On-Demand Scheduler 70, local matcher 72 and content delivery system (cds) server 74 components reside at the ISP POP server 16. The ISP POP

server 16 also includes a remote local database 76 storing individual user profile data and advertisements. The Dynamic Campaign Manager component 50 provides a portal to the system for advertisers (or Ad buyers or media buyers who act on behalf of advertisers) to initiate and manage their advertising campaigns. (The  
5 terms Ad buyer, media buyer, and advertiser are used interchangeably for purposes of this application.) The advertiser can, e.g., monitor the number of times content has been delivered to a client 10 (i.e., the number of impressions) and the number of click-throughs on that content during the course of a campaign.

The Capacity Forecaster 52 reviews new campaigns proposed by advertisers  
10 and predicts whether their campaign objectives are achievable in view of forecasted inventory of user screen real estate. The Capacity Forecaster 52 thereby assists in forming contracts having an expected high degree of success. A 'contract' as used herein is generally an agreement for content delivery typically between the scheduler system operator or owner and an advertiser or media buyer. This  
15 agreement specifies various terms including, e.g., the content to be delivered, delivery quantity (i.e., number of impressions), target subscriber group, and start and end dates.

The Inventory Manager 51 generates a candidate plan to fulfill new and existing advertiser contracts and to optimize usage of surplus user screen real  
20 estate. The Inventory Manager 51 modifies the plan as needed based on delivery feedback information received from the On-Demand Scheduler 70. The Delivery Manager 54 generates the best plan for the current days scheduled deliveries. The Delivery Manager 54 is also responsible for balancing deliveries within a day and processing the results of the On-Demand Scheduler. The On-Demand Scheduler 70  
25 dynamically constructs delivery schedules for individual users on user login. The Capacity Forecaster 52, Delivery Manager 54 and On-Demand Scheduler 70



components of the system will be described in greater detail below.

FIGURE 3 is a flowchart generally illustrating the process of scheduling and delivering content in accordance with the invention. At Step 100, the Dynamic Campaign Manager 50 receives a proposed new advertising campaign from an advertiser. At Step 102, the Capacity Forecaster reviews the proposed campaign to determine whether the campaign goals are achievable. If the inventory projections are less than the campaign goals, then at Step 104, the scheduler system identifies and suggests which constraints could be relaxed in order to achieve campaign goals as will be described below. If the campaign is determined to be achievable, it is approved at Step 106.

The Inventory Manager 51 then constructs a delivery plan for all approved campaigns at Step 108. At Step 110, for each subscriber on user login, the On-Demand scheduler 70 constructs an individual delivery schedule. At Step 112 advertisements are transmitted to users based on their individual delivery schedules. The system reports advertisement delivery feedback to the Delivery Manager at Step 114, which is used to update the master delivery plan as will be discussed below.

#### Capacity Forecaster

The Capacity Forecaster component 52 assists in predicting the success of a campaign proposed by an advertiser. For example, it predicts whether the system will be able to deliver a proposed number of impressions to users of some given profile within a desired period of time. When a new contract with an advertiser is under consideration (i.e., the advertiser proposes a new campaign), the Capacity Forecaster 52 calculates the probable or expected supply (i.e., surplus) of screen real estate on user client devices and approves the contract if an adequate supply is

expected for the proposed campaign. If the supply is not sufficiently large, the Capacity Forecaster 52 assists the advertiser in modifying the campaign requirements or constraints set by the Ad buyer by determining which constraints could be modified and how in order to successfully schedule a potential contract.

5           The Capacity Forecaster 52 determines campaign achievability by examining the number of qualified subscribers who match the campaign's profile using the Matcher and then calculating the current load (or alternatively, the available capacity, i.e., surplus) on those users due to previously scheduled campaigns. More particularly, the Capacity Forecaster 52 calculates the current load (or surplus)  
10       based upon statistics of how many impressions a subscriber can support compared to how many advertisements the subscriber is already scheduled to receive due to existing campaigns. (The statistics include data previously gathered on subscribers such as, e.g., how frequently a given subscriber logs on and the lengths of his or her sessions.) The total load on each subscriber is calculated by summing the loads  
15       contributed by each campaign. If a sufficient number of available subscribers is not available or if the load on the subscribers is too high, the Capacity Forecaster 52 identifies and suggests constraints to relax so that the campaign goals can be met such as, e.g., increasing the campaign length, reducing the number of requested impressions, or relaxing the profile constraints.

20           Additionally, the Capacity Forecaster 52 can periodically re-evaluate campaigns currently under execution, and determine their probability of success, e.g., whether the system will be able to schedule the contracted number of content deliveries based on delivery data feedback that has been received. Again, the Forecaster 52 can determine whether constraints set by the advertiser should be  
25       relaxed in order to increase the likelihood of success of the campaign.

#### Inventory Manager

The Inventory Manager 51 generates a master delivery plan expected to fulfill delivery contracts with advertisers. It uses delivery feedback information received from the On-Demand Scheduler 70 of each POP server 16 in the system to adaptively modify the master plan on a periodic basis.

5 For each accepted advertising campaign, the Inventory Manager 51 calculates a daily goal number of impressions to meet contract requirements. Advertisers typically desire to distribute the total number of desired impressions equally over each day of the campaign. (Alternatively, other distribution patterns can be used as desired.) The goals are periodically updated, e.g., each day, by comparing the  
10 number actually delivered to the desired total number of impressions.

The Inventory Manager constructs the master delivery plan on a periodic basis, e.g., once a day, based on the calculated goals of each of the active advertising campaigns. The plan specifies a prioritized master list of advertisements, which is sent to the On-Demand Scheduler 70 at each POP server 16. The prioritized content  
15 list identifies the order in which advertisements are to be displayed. The order is based preferably both upon priority and some weighting mechanism that indicates how many impressions are needed by each campaign.

The Delivery Manager 54 can reorder or reprioritize the master list of scheduled advertisements based upon delivery feedback data and queuing  
20 logic/algorithms. For example, if the goal for a given campaign is to evenly distribute an advertisement over the course of the campaign length, the advertisement can be moved down in the queue of advertisements to be displayed if it gets ahead of its daily goals. Similarly, if an advertisement gets behind in meeting its goals, it may be automatically promoted in priority. If an advertisement  
25 exceeds its daily goal it can be effectively shut off by being placed at the very end of the queue.

On-Demand Scheduler

The On-Demand Scheduler 70, which resides at each ISP POP server 16 in the system, dynamically constructs an individual ordered list of advertisements to be delivered for each given user upon user login. Each individual list includes  
5 advertisements matched to the user and prioritized according to the master list received from the Delivery Manager 54.

When a user or subscriber logs on, the On-Demand Scheduler 70 calls the Local Matcher 72 to identify the advertisements the subscriber is eligible for. This is done by matching the profile of the user (preferably through an anonymous user  
10 ID) to the profiles of advertisements in the prioritized master list generated by the Delivery Manager 54. The result is a list of advertisements for which the subscriber is eligible to receive. The On-Demand Scheduler 70 then preferably retrieves any state variables for the advertisements. The variables can indicate, e.g., if a  
15 subscriber has previously been sent an advertisement, when (date and time) that subscriber was last sent the advertisement, and how many total times the subscriber has been sent the advertisement.

After the On-Demand Scheduler 70 has the list of eligible advertisements and the state variables for those advertisements, it preferably examines any constraints for each advertisement. Constraints can include, e.g., the time required  
20 between successive impressions of a given advertisement, maximum number of impressions allowed, etc. If the constraints do not rule out the advertisement, it is placed in the subscribers' individual list of advertisements. The individual list is prioritized in accordance with the prioritized master list generated by the Delivery Manager 54.

25 The system employs on-demand scheduling for individual users to reduce

the communication volume over a network. On-demand scheduling avoids the need for replicating the large client command tables over the Internet from the master server 18 to the POP servers 16.

Advertisements in the individual list are retrieved by the CDS server 74 from the remote database 76 and sent to the subscriber in the specified order. The On-Demand Scheduler 70 preferably provides the Delivery Manager 54 with data on the impressions delivered to each user so that the Delivery Manager 54 can reprioritize its list of advertisements as needed.

#### Scheduling Mathematics

The Capacity Forecaster 52 and the Inventory Manager 51 determine campaign achievability and construct the master delivery plan based on statistical behavior of subscribers (obtained, e.g., as described in the above-mentioned application Serial No. 09/558,755), generally taking into consideration a number of random factors. One random factor, e.g., is that the system does not know whether a given user will log on at any given time period. Another random factor is that the amount of time the subscriber will be logged on (i.e., the length of an online session) is unknown. Accordingly, the number of advertisements that can be delivered to a particular subscriber for some given time period is unknown. The system models these random factors and uses probability distributions, expected values and related statistical computations based on knowledge of prior observed online behavior of users.

These random factors can be mathematically represented, e.g., as follows. A random subscriber will log onto the Web on any given day with a probability  $p$ . He or she will therefore not log onto the Web on any given day with probability  $(1-p)$ .

Given that a subscriber logs onto the Web on a given day, suppose that a random

subscriber's probability distribution function for Web time per day is given by

$F_s(s) \equiv P\{S \leq s\}$  = Probability that the service time is less than or equal to a given value, say  $s$ ,

- 5 where  $S$  is a random variable indicating the subscriber's Web time per day (given that  $S$  is positive).

Suppose advertisements are to be sent to the subscriber once each  $T$  minutes for up to  $A_{max}$  ads (e.g.,  $A_{max} = 10$  ads per day). That is, the first advertisement is sent  
 10 at time  $T$  minutes, the second at time  $2T$  minutes and the  $k^{th}$  at time  $kT$  minutes,  $k \leq A_{max}$ .  $[1 - F_s(s)]$  is the probability that the Web subscriber remains on the Web on a random day for more than  $s$  minutes. Suppose that there are  $A_{max}$  ads that the system would like the subscriber to see on any given random day. The probability that the system would be successful at delivering the  $k^{th}$  such advertisement is

15 
$$P\{k^{th} \text{ ad delivered}\} = p[1 - F_s(kT)] \equiv p_k, \quad k = 1, 2, \dots \quad (1)$$

Equation (1) indicates that the probability that the  $k^{th}$  advertisement slated for delivery will be successfully placed is equal to the product of two quantities: the probability  $p$  that the subscriber logs onto the Web on a given day and the  
 20 conditional probability (given that he or she has logged on) that he or she remains logged on for at least  $kT$  minutes. Since  $F_s(s)$  is a monotone non-decreasing function of  $s$ ,  $p_k$  is a non-increasing function of  $k$ . In other words, the likelihood of successfully placing an advertisement becomes less and less as the advertisement is placed lower in the advertisement priority list of the master plan generated by the  
 25 Inventory Manager 51.

The mean number of advertisements placed to a random subscriber on a

random day can be represented as follows:

$$\bar{A} = \sum_{k=1}^{A_{\max}} p[1 - F_S(kT)] = \sum_{k=1}^{A_{\max}} p_k. \quad (2)$$

5 The probability mass function for the number of advertisements placed can be found by taking successive differences in the cumulative distribution function for the random variable  $S$ , i.e.,

$$\begin{aligned} P\{A = k\} &= P\{S < (k+1)T\} - P\{S < kT\}, \quad k = 0, 1, 2, \dots \\ P\{A = k\} &= F_S([k+1]T) - F_S(kT) \\ P\{A = A_{\max}\} &= 1 - P\{S < A_{\max}T\} = 1 - F_S(A_{\max}T) \end{aligned} \quad (3)$$

10 The second moment of the number of advertisements placed can be represented as follows:

$$\overline{A^2} = \sum_{k=1}^{A_{\max}-1} k^2 \{F_S([k+1]T) - F_S(kT)\} + A_{\max}^2 [1 - F_S(A_{\max}T)] \quad (4)$$

15 An expression for the variance of the number of advertisements placed on a random day to a random subscriber can be written as follows:

$$\text{VAR}(A) \equiv E[(A - \bar{A})^2] \equiv \sigma_A^2 = \overline{A^2} - \bar{A}^2 \quad (5)$$

20 where the formulas for  $\bar{A}^2$  and  $\overline{A^2}$  are from Eqs.(2) and (4), respectively, and where  $E[*]$  is the expectation operator, which is defined further below with respect to Eqn(13).

Given the above foundational elements, an engineering approximation of

queue ordering can now be calculated.

Suppose  $E[Q_j] = \text{"average Queue placement for ad type } j \text{ on day } d\text{"}$  is thought of as an engineering approximation to its distribution over Queue placements. For instance, if  $E[Q_j] = 11.4$ , then ad type  $j$  averages between the 11<sup>th</sup> and 12<sup>th</sup> place in all individual ordered lists that include ad type  $j$ . Sometimes ad type  $j$  may find itself in a front position in an individual ordered list, perhaps even first or second; sometimes it may trail far back in the pack, perhaps at position 30 or so. The number 11.4 represents an average position over all individual ordered lists that are generated and that include ad type  $j$ .

Derived here an inequality that provides useful information on the required individual ordered list priority of a given ad.

Use the definition:

$M_j = \text{average number of ads to be placed per day}$  (For simplicity this ignores complications like holidays and weekends.)

Also we call  $P_j$  the population of subscribers who are eligible for ad type  $j$  and we recall that  $p$  is the fraction that log on during a random day. Thus the product  $Pp$  is the expected number of subscribers who log on during a random day and who are eligible for ad type  $j$ .  $F(x)$  is the probability that a logged on subscriber will have a session time less than or equal to  $x$  minutes in length. This means that  $F(\Delta E[Q_j])$  is the probability that a logged on subscriber will have a session time less than or equal to  $\Delta E[Q_j]$  minutes. Consequently,  $\{1 - F(\Delta E[Q_j])\}$  is the probability that a logged on subscriber will have a session time *greater than* to  $\Delta E[Q_j]$  minutes. Thus, recognizing the LHS (left hand side) as the expected number of successful ad  $j$  placements during



a day, we can write:

$$P_j \{p(1 - F(\Delta E[Q_j])) \geq M_j\}$$

5 This equation can be manipulated to obtain:

$$1 - F(\Delta E[Q_j]) \geq M_j / (P_j p) = \text{required number of ads to be delivered divided by the expected number of eligible users who log on}$$

(12)

10

This provides an upper bound of sorts on the maximum feasible value for  $E[Q_j]$ . That is, for a given parameter set, ad  $j$  should be placed at least at the position in the queue ("as experienced" by ad  $j$ ) indicated by this inequality. Every night, the average value of  $E[Q_j]$  can be found for each ad by using Forecaster, sampling

15 without replacement for the entire population of subscribers, and tabulating separately for each on-going campaign.

20

The following is an example of a case in which the subscriber holding time is a negative exponential random variable, with mean  $\tau$ . In recent statistical work, it has been found that the distribution of Web time per logon is negative exponential and its mean  $\tau$  is approximately 31 minutes. In the case of a negative exponential distribution, we can write

$$F(x) = \Pr\{\text{session time} \leq x\} = 1 - e^{-x/\tau}$$

25

Thus Eq.(12) can be written in this case as

$$e^{-\Delta E[Q_j]/\tau} \geq M_j / (P_j p)$$

The natural logarithms of both sides can be taken to obtain

$$-\Delta E[Q_j]/\tau \geq \ln(M_j / (P_j p)) = \ln M_j - \ln P_j - \ln p.$$

(The natural logarithm is a strictly monotonic transformation, thereby preserving inequality directions. For example if  $e^x$  is greater than or equal to some quantity, then  $x$  is greater than or equal to the natural logarithm of that quantity.)

Manipulating this equation, the desired constraint on the expected queue position can be obtained:

$$E[Q_j] \leq -\tau(\ln M_j - \ln P_j - \ln p) / \Delta \quad (13)$$

Consider the following numerical example: Suppose  $\tau = 31$  minutes,  $M_j = 10,000$  ads per day,  $P_j = 100,000$  subscribers who are eligible for ad  $j$ , and  $p = 0.15$  (probability that a random eligible subscriber logs on during a random day), and  $\Delta = 2$  minutes (time between ad impressions). Then, substituting into Eq.(13), we have

$$E[Q_j] \leq -\tau(\ln M_j - \ln P_j - \ln p) / \Delta = -31(\ln 10,000 - \ln 100,000 - \ln 0.15) / 2$$

or

$$\begin{aligned} E[Q_j] &\leq -31(\ln 10,000 - \ln 100,000 - \ln 0.15) / 2 = -31(9.21 - 11.51 - (-1.90)) / 2 \\ &= -31(-0.4) / 2 = 6.2. \end{aligned}$$

This means that in order to deliver the required number of type  $j$  ads today, the average individual ordered list position for type  $j$  ads must be 6.2 or less. (As used herein, lower position numbers represent higher priorities). A sensitivity analysis can be easily be performed here. Suppose the inter-ad placement time  $\Delta$  is

5 doubled to 4 minutes. Then  $E[Q_j]$  must be 3.1 or less, a much more stringent constraint due to fewer total ads being placed. But suppose that  $\Delta$  remains fixed at 2 minutes and that the mean time per session logon increases 50% to 46.5 minutes.

Then  $E[Q_j]$  can increase up to 50% above its previous value, to 9.3. So it can be seen how the individual ordered list inequality yields constraint values that are

10 directly proportional to the mean time per session  $\tau$  and inversely proportional to the time between ads placement  $\Delta$ . The functional behavior with respect to population count, number of ads to be placed and probability of logon all are more complicated due to the natural logarithms. But a key observation is that this equation (13) can readily be used to perform sensitivity analyses on any combination of these

15 quantities.

Additional mathematics may be employed to embody the optimization of the internal master list ordering. A so-called Delivery Specification Optimizer (DSO) software module can embody the mathematical properties to adjust the ordering of the queue list. Specifically, the DSO attempts to find a queue list  $S$  such that, for each

20 ad  $A$ , the "delivery ratio":

$$\text{DeliveryRatio}(A) = E(\text{number of deliveries for } A \mid S) / \text{goal}(A),$$

where  $E(\text{number of deliveries of } A \mid S)$  is probabilistic shorthand for number

25 of impressions of  $A$  the master list ordering  $S$  will deliver on an average day is the same for all advertisements.

The delivery ratio is used to assess the performance of an individual content

delivery in proportion to the master list. This provides the basis of evaluating the efficiency of the master list ordering. Different queue list orderings are evaluated according to the DeliveriesExpected they yield compared to the input DeliveriesRequested. The DSO attempts to achieve a constant value of

- 5 DeliveriesExpected / DeliveriesRequested over all ads; or as close to this optimum as is possible. In other words, the algorithm attempts to minimize the fitness value, where:

The fitness value = the sum of squares of ( DeliveryRatio(ad) – DeliveryRatio( { all ads} ) ) (Eqn. 14)

10

The DSO could optionally utilize numerical analysis methods to derive the relative importance of each piece of content in the queue list. The over all fitness of a queue ordering is the sum of the values of all of the contracted content deliveries. This value is a differentiable function with respect to the weightings. As a result, the

15 DSO can use multi-dimensional Newton's method or the steepest descent method to rapidly approach an optimal solution.

#### Delivery Qualities

- In accordance with a further embodiment of the invention, the scheduler system preferably ensures that subscribers always have content to display even if
- 20 they are not eligible for any active campaigns. Accordingly, the system preferably provides a set of default or filler impressions to be displayed when there is no content available for a given user.

In accordance with a further embodiment of the invention, the scheduler system is capable of delivering 'instant' advertisements (or other content) to subscribers. These are advertisements that are delivered to users if they perform some given action, e.g., enter a particular keyword in a search engine query or enter a particular URL address. The system can preferably preempt the normal queue of

ordered advertisements in an individual schedule with an instant advertisement when needed. The system preferably allows the percentage of time that instant advertisements can preempt the normal queue to be configurable in order to reduce errors in calculations made by the Capacity Forecaster.

The system improves use of excess inventory. It can also increase the likelihood of over-delivery (i.e., delivering a greater number of impressions than requested by an advertiser), which is typically favorable to advertisers. It provides a generally even distribution of impressions over the length of the campaign (if so  
5 desired). The system provides greater diversification of impressions (i.e., the advertisements are distributed to different users in a target group).

Having described preferred embodiments of the present invention, it should be apparent that modifications can be made without departing from the spirit and scope of the invention.